REMARKS

Claims 1-4, 6-9, 11-13, 16-17, 19-22, 32, 34-35, 38, 40-41, 44, 46-47, 50, 52-53, 56, and 59-64 remain pending in the application. Claims 1, 3, 7, 9, 11, 17, 20, 34, 35, 38, 40, 41, 44, 46, 47, 50, 52, 53, 56, 60, 61, 63, and 64 have been amended without introduction of new matter. Favorable reconsideration is respectfully requested in view of the above amendments and the following remarks.

Before addressing the issues raised in the Office Action, it has recently come to Applicants' attention that the filed Declaration is defective because it doe not identify the citizenship of inventor Jeff BOND. In response, a new declaration of inventor Jeff Bond is submitted herewith that complies with 37 CFR §1.67(a).

Regarding the Specification, the title of the invention was objected to as allegedly not being descriptive. In response, the title has been amended to read "Data Processing Architectures for Packet Handling". As this is believed to address the Office's concern, withdrawal of the objection is respectfully requested.

The Office made a number of objections to the claims. These are traversed as follows.

In one aspect, claims 11, 34, 40, 46, and 52 were objected to "because the claims do not refer back to another claim as mandated by 37 CFR 1.75(c), as they are all dependent on claim 60." This objection is believed to be improper. MPEP §608.01(o) IV page 600-92 (Rev. 7, July 2008) instructs: "During prosecution, the order of claims may change and be in conflict with the requirement that dependent claims refer to a preceding claim. Accordingly, the number of dependent claims and the numbers of preceding claims referred to in dependent claims should be carefully checked when claims are renumbered upon allowance."

The MPEP 608.01(n) F page 600-88 further instructs: "However, in situations where a claim refers to a numerically following claim and the dependency is clear, both as presented and as it will be renumbered at issue, all claims should be examined on the merits and no objection as to form need be made. In such cases, the examiner will renumber the claims into proper order at the time the application is allowed."

The facts concerning the present claims match those described above in the MPEP. Each of claims 11, 34, 40, 46, and 52 originally referred to claim 10 (i.e., a previous claim). However, the order of claims changed during prosecution: Claim 10 was cancelled and claim 60 was added in the Amendment filed on April 18, 2005. The cancellation of claim 10 required that the dependencies of claims 11, 34, 40, 46, and 52 be changed, and it was desired

that they depend from the new claim 60. In each case, the chain of dependencies is clear, as claim 60 depends from claim 9, which in turn depends from claim 1.

Notwithstanding the above, the issue has been rendered moot because claim 11 has been amended to include the input/output systems, as per claim 60, and has been made to refer back to claim 9. Furthermore, each of claims 34, 40, 46 and 52 has been amended to depend from claim 1.

For at least the foregoing reasons, the objection to claim 11, 34, 40, 46, and 52 should be withdrawn.

In another aspect, claim 60 was objected to because of the recitation "input output system." In response, this has been amended to now recite "input/output system", as suggested by the Office.

In another aspect, claim 17 was objected to because of the recitation "data I/O". In response, claim 17 has been amended to expressly recite "data input/output system", as suggested by the Office.

Finally, claim 20 was objected to because of the recitation "SIMD". In response, claim 20 has been amended, as suggested by the Office, to expressly define the acronym as follows: "a single instruction multiple data (SIMD) ...."

For at least the foregoing reasons, it is respectfully requested that the objections to claims 60, 17, and 20 be withdrawn.

Claims 1-4, 6-9, 11-13, 16-17, 19-22, 32, 34-35, 38, 40-41, 44, 46-47, 50, 52-53, 56, and 59-64 stand rejected under 35 U.S.C. §112, second paragraph, as allegedly being indefinite. This rejection is respectfully traversed.

As regards claim 1, the Office expressed the concern that it is indefinite whether the expressions "data packet" and "whole data packet" refer to two different things. In order to remove any scope for doubt, all occurrences of the word "whole" have been removed from claims 1, 3, and 61.

As regards claim 3, the Office objected that the word "it" in the third line should be explicitly recited. In response, the word "it" has been replaced by "said packet" in order to make the meaning abundantly clear.

As regards claim 7, the Office objected that the claim is not clear as to the nature of the "control" exerted on the input device by the PEs, suggesting instead that any such control is exerted by the thread sequence controller and not by the PEs. In reply, reference is made to the specification at page 25, lines 5-6, which states that the processors always initiate

datagram transfers within the system domain. Referring more specifically to the paragraphs spanning page 52, lines 3-18 of the specification, for example, the processor of the particular embodiment described operates by multi-threading. Each processor runs the same program, which contains four threads, identified as loading, processing, lookup and unloading. Bearing in mind that the PEs initiate datagram transfer, this means that when a PE is ready to perform an operation, it signals its readiness to the processor to switch to a packet loading thread, whereby to cause the thread to load new packet data into its memory for processing by the packet processing thread. The distributor forwards (batches of) packets to a processor in response to the request from the PEs (see, e.g., the specification beginning at page 26, line 6). Moreover, as stated in the paragraph spanning page 26, lines 12-20, the processor requests data to be sent from a data buffer in chunks (i.e., packet portions) whose size and number are specified by the processor. This paragraph emphasizes that "the processor retains control of the transfer". In this sense, the PEs control the allocation of packet portions to themselves, the size of the packet portions and the number of packet portions. These paragraphs therefore support the concept that data transfer is initiated in response to a request from the processor, as claimed in claim 7. Claim 7 as amended now makes this clear by stating that the processing elements are operable to control the transfer of packet portions to the processing elements from the input device. Claim 8, which depends from claim 7, has therefore been similarly clarified.

As regards claim 9, the Office objected that the claim is not clear as to the nature of the "control" exerted on the output device by the PEs. In response, by analogy with claim 7, the PEs control the output of data packet portions from the processing elements to an output device after processing in the PEs. A PE that has completed an operation signals the thread scheduler (via the processor) to switch threads to unload the processed data to the output device. In this sense, the PEs control the output of data packets from the PEs. Claim 9 as amended now makes this clear by stating that the processing elements are operable to control the output of data packet portions from the processing elements to an output device.

As regards the numerous instances quoted by the Office in which there is an alleged inconsistency as to the nomenclature of specific features in some of the dependent claims, the Office's attention is respectfully directed to the following amendments that are believed to address the Office's concerns:

- claim 17 now refers to a data "input/output system" rather than an "I/O structure";

- claims 34-35, 38, 40, 44, 46-47, 50, 52-53 and 56 now refer to "A data processing architecture" rather than merely to "An architecture"; and
- claims 52-53 no longer include the words "or system", to remove the lack of clarity objected to by the Office.

Claims 63-64 are objected to because of the limitations relating to "the bandwidth" and "the amount of required processing". In response, each of claims 63 and 64 has been amended to now recite "a bandwidth" and "an amount of required processing." Claim 63 requires the number of processing elements (PEs) to be determined based on the bandwidth and the amount of required processing. The references to "bandwidth" and "amount of required processing" are supported in the specification at, for example, the paragraphs spanning page 55, line 28, through page 57, line 3, which discuss how the number of PEs can be determined. The paragraph beginning at page 55, line 28, starts with the premise that a desired level of performance can be provided by choosing an appropriate number (n) of PEs in relation to the bandwidth of the data stream and the minimum size of packets, leading to figure of 104 million packets per second. This may be expanded by the factors discussed in the paragraphs spanning page 56, line 4 through page 57, line 3, which relate to processing capability. Additional support can be found in the paragraphs spanning page 7, line 25 through page 8, line 5 of the specification, which emphasis the "bandwidth-centric" approach of the invention. In addition, the paragraphs spanning page 7, line 25 through page 8, line 1, specifically support the feature that "processing and storage should be allocated by bandwidth, not by packet....".

As regards claim 64, the paragraph spanning page 26, lines 12-20 explicitly states that the processor, comprising the PE array, requests transfer of data from a buffer, wherein the request specifies the number of chunks (i.e. ,data portions) and the maximum chunk size. The paragraphs spanning page 7, line 25 through page 8, line 5 of the specification refer to processing being allocated by bandwidth. Also, original claim 1 refers to the distribution of data packets being dependent on the data processing bandwidth of the PEs. The specification paragraphs spanning page 62, line 4, through page 63, line 25 discuss the scalability of the platform, including adjusting the number of processors to achieve higher performance; that is, processor numbers as a function of performance, the number of processors in a cluster, and functionality versus performance defined in software. Such statements, taken together, in the context of the specific implementations presented in the described embodiments, support the

feature of claim 64 that the size of the packet portions is determined based on the bandwidth and the amount of required processing.

For at least the foregoing reasons, claims 1-4, 6-9, 11-13, 16-17, 19-22, 32, 34-35, 38, 40-41, 44, 46-47, 50, 52-53, 56, and 59-64 are believed to define the invention with sufficient particularity and distinctness to satisfy the requirements of the statute. It is therefore respectfully requested that the rejection of these claims under 35 U.S.C. §112, second paragraph, be withdrawn.

Claims 1-4, 6-9, 11, 16-17, 20-21, 32, 40, 44, 52, 56, 59-60, and 63-64 stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Marsan et al., "Router Architectures Exploiting Input-Queued, Cell-Based Switching Fabrics (henceforth, "Marsan") in view of Clauberg (WO 97//29613 -- henceforth "Clauberg"). This rejection is respectfully traversed.

The Marsan paper discloses in its Abstract and in the passage in page 2 indicated by the Office, the concept of variable-sized incoming packets being segmented into ATM-like cells; that is, fixed-size cells. Trains of these cells are then processed internally within a switch fabric and are then reassembled to form outgoing packets on an output path. The cells in any one train are not interleaved with the cells of another train, so the cells of a train remain contiguous within the switch fabric. So, Marsan teaches that an incoming stream of variable-sized packets can be chopped up into fixed-size portions (cells) analogous to the cells of an ATM system.

As a matter of no small significance, Applicants' method does not require incoming packet portions to be divided into fixed-size portions. Indeed, a cursory glance at Figure 2 of the instant specification illustrates graphically that the size of the data portions placed in PE0 to PE12 are anything but equal or fixed in size. The paragraph spanning page 26, lines 12-20 explicitly states that the processor requests data in chunks, whose number and maximum size are specified by the processor.

This is one of the benefits that derive from embodiments of Applicants' invention, in that the processor can simply take anything that is "thrown" at it and can "chop" it into manageable portions that are then distributed over any number of PEs. Processing of these portions is then carried out and the processed portions "thrown" out to an output path for further processing. This rather informal language is used deliberately to emphasize just how robust, adaptable, versatile, powerful and capable Applicants' processor architecture actually is.

Further, in Marsan, the whole object is to include a switch fabric within the packet processing operation. In stark contrast, at no time is any switching carried out in Applicants' method and apparatus. This makes the whole operation quicker and simpler. The significance of this is that the person of ordinary skill in the art is less likely to refer to Marsan as a starting point for a solution to the problem of handling packets of unpredictable size because Marsan is intimately tied up with a switching problem where different constraints apply.

Turning now to the Clauberg published patent application, Clauberg discloses parallel processing of fixed length cells. The term "parallel processing" is not used by Clauberg in the usual sense commonly employed in parallel processors such as SIMD processors. Rather, Clauberg uses the term in the sense that an incoming stream of fixed sized cells is split into a number of parallel paths for processing in parallel to one another. Processing in each path, however, is strictly sequential. As shown in Figure 1 of Clauberg, a stream of fixed size input cells from an ATM "time-slotted" input path 10 is sent one cell at a time to a respective one of a number of parallel paths 13.1 to 13.5. At any instant, there are only two such cells being processed by the first or the sequential second processor 14.1 or 15.1, etc., in each path.

The cells are clocked through the respective paths at the same rate as the time slots of the input line, as illustrated in Clauberg's Figure 2. Measures are taken to ensure that a cell does not reach the second processor of a path before the first processor on that path has completed its process on that cell (see page 9, of Clauberg, for example). This is to ensure that the cell order is preserved as the cell sequence 10 passes through the switch fabric 9.

The Office's attention is respectfully directed to the following differences between Clauberg and Applicants' claimed invention. In Clauberg, a strict succession is adhered to, where cells arriving at an input multiplexer 12 are sent to successive paths 13.1-13.5 in a specified order. The cells are effectively sent to the respective paths in a round-robin sequence. There is therefore no "distribution" as such in the sense that the term is used in Applicants' invention; that is, there is no question of cells being sent to processing elements (equivalent to the processor units of Clauberg) as and when those elements become free to take data for a further processing operation. Rather, the cells in Clauberg have the same destiny as all the others.

In Clauberg, there is a considerable amount of idle time in the processors 14.1, 15.1, etc., because of the round-robin nature of cell transfer to the parallel paths between input and output. In true parallel processors, such as in the instant invention, efforts are made to ensure

that as many PEs as possible are active for as much time as possible. In Clauberg, the architecture makes it inevitable that most of the PEs are idle for most of the time. In Applicants' processor, packet portions are distributed to PEs that have completed a processing task and are ready to receive new data. Batches of data packets or packet portions are distributed over available PEs wherever those PEs are located in the array. There is no necessity for a strict allocation as in Clauberg.

Clauberg does not actually distribute packets (or cells) in this sense at all. The Clauberg switch fabric is robotic in that incoming cells are sent to respective paths in a predetermined order. There is no element of the cells being distributed to the paths in dependence on the size of the incoming cells because the cells are all the same size. There is certainly no "dynamic" distribution, as required by Applicants' claims.

Moreover, and this is fundamental to the rejection, even if one were to accept that the features of Applicants' claims would be obvious on the basis of a combination of the teaching of Marsan and Clauberg, which Applicants expressly do not accept, there is a feature in Applicants' claims whose significance has not apparently been appreciated, which will now be addressed.

The last two lines of claim 1 state, "wherein the data processing architecture is operable to process at least one data packet at a time". If one were to try to relate this feature to the Marsan/Clauberg combination, Marsan takes incoming packets and chops them into cells and then Clauberg "distributes" these cells to the parallel paths of its switch fabric. However, as discussed above, Clauberg is only capable of processing two cells at any one time. Clauberg is therefore incapable of processing at least one packet at a time. Within this limitation, it may be noted that there can only be one cell at a time in any "column" of processing units in the Clauberg architecture. The overall combination is therefore restricted by this limitation, which is inherent in Clauberg. Consequently, the quoted feature of Applicants' claim cannot possibly be met by any combination of Marsan with Clauberg. The only exception to this is the hypothetical situation in which the original packet is of such a size that it can be chopped by Marsan exactly into two ATM-like cells, which would then be placed in succession on two of Clauberg's paths. However, it is respectfully submitted that such an instance would be a remarkable coincidence, which does not seriously challenge the inventiveness of Applicants' processor architecture. The paragraph spanning page 8, lines 18-23 of Applicants' application is one instance where this feature is specifically supported in the description - "Entire packets are loaded into the processors".

Even though Applicants are firmly of the view that their claims as now presented are clearly differentiated over Marsan in view of Clauberg, it is also desired to bring to the Office's attention the fact that claim 1 as amended also includes a further limitation that the input device is operable to distribute the data packets in batches across the processing elements. Support for this limitation can be found in various paragraphs throughout the specification but particular attention is directed to the paragraphs spanning page 26, lines 6-11 and page 27, lines 19-30, which refer to the processor requesting datagrams (packets) and "chunks" (packet portions) using a "batch read" mode; and the paragraph spanning page 59, lines 15-20, which states categorically that each processor handles a batch of packets sufficient to fill the local memories of its PEs. Further references occur in the paragraph spanning page 42, lines 24-30, which introduces the data buffer as instrumental in providing batch transfer; the paragraph spanning page 53, lines 8-17, which highlights concurrency of various operations including batch table lookups; the paragraph spanning page 54, line 27 through page 55, line 9, which explains the application of multiple passes and the relationship with packet input to the processor being in batches; and the paragraph spanning page 56, lines 4-14, which gives specific examples for the number of packets within a batch, viz 64 packets, each of 40 bytes.

The significance of the packets being distributed in batches underlines the previous arguments as to the manner in which incoming data packets are distributed to the PEs and how the inventive method of the instant invention differs significantly from the prior art. This difference is also reinforced in the fact that the person of ordinary skill in the art would not have arrived at the combination of Marsan and Clauberg if he had been seeking a solution to the unpredictable size data packet handling problem addressed by the present invention. Even if he had, and without the benefit of hindsight, the combination does not, in any case, lead to the instant invention as currently claimed because of the ability of the instant invention to process at least one data packet at a time. If that were not enough, the fact that the processor as implemented operates through multithreading, where the PEs are programmable (see the paragraphs at page 60, line 9 and page 63, line 27 through page 64, line 2, for example), takes Applicants' invention to a different level of versatility and utility as compared to the robotic and limited switch fabric-type applications exhibited in the prior art Marsan and Clauberg.

For at least the foregoing reasons, it is respectfully contended that the subject matter as now claimed is neither disclosed nor suggested by the prior art of record. On that basis,

any claims that are dependent from independent claim 1 are new and non-obvious for at least the foregoing reasons, so the objections to their patentability also become moot. It is therefore respectfully requested that the rejection of claims 1-4, 6-9, 11, 16-17, 20-21, 32, 40, 44, 52, 56, 59-60, and 63-64 under 35 U.S.C. §103(a) be withdrawn.

Claims 12-13, 19, 34-35, 38, 41, 53, and 61-62 stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Marsan and Clauberg as applied to claims 1, 2, 17, and 60 above, and further in view of Kejriwal et al. (US 6,704,794 -- henceforth "Kejriwal"). This rejection is respectfully traversed.

Claims 12-13, 19, 34-35, 38, 41, 53, and 61-62 all depend from independent claim 1, and are therefore patentably distinguishable over the Marsan and Clauberg patents for at least the reasons set forth above. The Kejriwal patent fails to make up for the deficiencies outlined above with respect to Marsan and Clauberg, so that any combination of Marsan, Clauberg, and Kejriwal would still lack at least the features discussed above with respect to claim 1.

For at least the foregoing reasons, it is respectfully requested that the rejection of claims 12-13, 19, 34-35, 38, 41, 53, and 61-62 under 35 U.S.C. §103(a) be withdrawn.

Claims 22, 46, and 50 stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Marsan and Clauberg as applied to claims 1, 17, and 60 above, and further in view of ISSC95 (Evening Discussion Session). This rejection is respectfully traversed.

Claims 22, 46, and 50 all depend from independent claim 1, and are therefore patentably distinguishable over the Marsan and Clauberg patents for at least the reasons set forth above. The ISSC95 document fails to make up for the deficiencies outlined above with respect to Marsan and Clauberg, so that any combination of Marsan, Clauberg, and ISSC95 would still lack at least the features discussed above with respect to claim 1.

For at least the foregoing reasons, it is respectfully requested that the rejection of claims 22, 46, and 50 under 35 U.S.C. §103(a) be withdrawn.

Claim 47 stands rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Marsan and Clauberg, and Kejriwal as applied to claim 12 above, and further in view of ISSC95. This rejection is respectfully traversed.

Claim 47 depends from claim 12, and is therefore patentably distinguishable over any combination of the Marsan, Clauberg, and Kejriwal documents for at least the reasons set forth above. The ISSC95 document fails to make up for the deficiencies outlined above with respect to Marsan, Clauberg, and Kejriwal, so that any combination of Marsan, Clauberg,

Kejriwal, and ISSC95 would still lack at least the features discussed above with respect to claim 1.

For at least the foregoing reasons, it is respectfully requested that the rejection of claim 47 under 35 U.S.C. §103(a) be withdrawn.

The application is believed to be in condition for allowance. Prompt notice of same is respectfully requested.

Respectfully submitted,

Potomac Patent Group PLLC

Date: June 25, 2009

By: /Kenneth B. Leffler, Reg. No. 36,075/
Kenneth B. Leffler
Registration No. 36,075

P.O. Box 270
Fredericksburg, Virginia 22404
703-718-8884